

Planning for Information Visualization in Mixed-Initiative Systems

Scott M. Brown

Crew Systems Interface Division
Air Force Research Laboratory
Wright-Patterson AFB, OH 45433
scott.brown@he.wpafb.af.mil

Michael T. Cox

Department of Computer Science and
Engineering
Wright State University
Dayton, OH 45435-0001
mcox@cs.wright.edu

Abstract

This paper describes two forms of information visualization for mixed-initiative systems associated with team collaboration and begins to discuss how plans might be formulated to achieve the visualizations. *Common understanding visualization* is concerned with visualizing the information a team employs, whereas *visual collaboration* is concerned with visualizing the ongoing, incremental information collection, the credibility and origins of that information, and the dynamic interpersonal relationships of the team itself. The first is the more “classic” form of visualization where data and information is collected, analyzed, abstracted, and tailored for display to the user. We are concerned not only with visualization for the single user, but also with visualizing the relationship the information holds in regard to the entire team. At the level of the individual user, a mixed-initiative system must consider how to tailor the appropriate information given the user’s skill, expertise, and preferences. At the corporate level, a system must manage the display of information across multiple human users and system components that share a common goal. The second form of visualization deals with the collaboration between a user and his/her mixed-initiative system, between users, and between systems. That is, the user needs to understand and visualize the collaboration, how the user fits into it, and the associated human/information interaction and processes involved in such collaboration. We claim both types of visualization are important for effective collaboration.

Introduction

The complexity of information intensive environments is overwhelming users. Mixed-initiative systems address this problem of increasing task load by providing assistance to users in information intensive environments. A system can help alleviate the user’s task load and increase situational awareness by collecting, extracting, and analyzing relevant information, by providing information abstractions of that information, and by helping the user manage the display of information at the user interface. Knowing the user and how to best present information to the user is central in this support.

The interaction between humans and information is a highly complex and dynamic task in which the acquisition, visualization, manipulation, and suppression of knowledge are all crucial supportive components to a user’s goals. In this respect, each above component of the interaction can be viewed as a mixed-initiative planning task from the perspective of an intelligent mixed-initiative system assigned to assist a human in an information environment. For example, wanting to acquire a piece of knowledge is equivalent to possessing a knowledge goal to change the current state of information. Such goals can be achieved by generating and executing knowledge plans that are analogous to plans to affect the state of the physical environment. Knowledge plans may involve actions such as information gathering activities on the internet and displaying graphical data at the user interface. Like standard plans these actions may possess constraints such as the restriction of particular classes of information.

One of the advantages of using a mixed-initiative system is the ability to allocate tasks, goals, and/or functions between a human user and the (computer) system to improve the individual performance (e.g., effectiveness and efficiency) of each. A user’s strength lies in the ability to provide guidance and insight concerning information that is necessary to draw complex, higher level inferences for data. The system’s strength lies in its ability to perform data acquisition and management, to include display of this information, from many heterogeneous sources, low level quantitative and qualitative analysis, and routine inference to enable decision support. A key to improving performance is allocating the tasks correctly and having a keen understanding of whom is doing what for/to whom.

In distributed environments where multiple users and their associated systems may be interacting collaboratively to achieve a goal, a “higher-level” knowledge plan emerges from the individual knowledge plans of each user. The ability for individuals to “see the big picture” of the emerging knowledge plan (e.g., who are the other users, what are the origins of the information, is the information credible, and how does one knowledge plan affect the overall plan) is vital to effective collaboration.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE JUL 1999		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Planning for Information Visualization in Mixed-Initiative Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Research Laboratory Wright-Patterson AFB, OH 45433-7022				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 9	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

This paper describes two forms of information visualization for mixed-initiative systems and speculates as to the manner for which the visualization can be planned. Specifically, we look at the collaboration of multiple users to achieve some goal and the planning for information visualization in mixed-initiative systems. We consider two types of visualization in distributed collaboration.

- *Common understanding visualization* is concerned with visualizing the information a team employs with respect to the team objectives but tailored for the individual team member.
- *Visual collaboration* is concerned with visualizing the ongoing, incremental information collection, the credibility and origins of that information, and the dynamic interpersonal relationships of the team itself.

We claim both types of visualization are important for effective collaboration. We investigate a number of critical research issues along these two directions. To highlight some of the issues, we discuss a simple version of planning a tailored information presentation and expand the scenario to the case of distributed collaboration. Subsequent sections will enumerate the components of visualization and expand on our claims.

Common Understanding Visualization

The main goal of information visualization can be summed up with the adage: “the right information to the right person at the right time.” The simplicity and succinctness of the goal belies the breadth and of the ways to achieve the goal. While there is much research on advanced information visualization, much of the research deals with single user/single system visualization. When multiple users must collaborate to achieve a set of goals, visualizing this “system-of-systems” becomes paramount. “Common understanding” does not mean “common visualization.” The ability of a system to effectively plan for information visualization for all users of a collaborating team relies on understanding the needs of the individual users, appropriately displaying the information based on these needs, and maintaining a consistent view of the emerging plan. This section describes some of the methods used to achieve the information visualization goal and some of the research issues associated with each method.

User Modeling

From studies of human interaction with complex systems in several domains, it is clear that sophisticated technology alone does not inherently increase system effectiveness (Mitchell & Sundstrom, 1997). Designing effective collaborative information in real-world applications is a complex endeavor. Although advances in computer technologies provide useful tools to support the organization and retrieval of information, purely technology-centered approaches often lead to problems in

user-system mismatches and result in inflexible and hard-to-use systems (Narayanan et al., 1997). Several factors, including the user's cognitive capabilities and limitations, the work domain and task constraints, the content and form of support from computer agents, influence the interactive collaborative planning process.

Providing intelligent assistance and performing tasks on the user's behalf requires an understanding of the goals the user is performing, the motivation for pursuing those goals, and the actions that can be taken to achieve those goals. User intent denotes the actions a user intends to perform in pursuit of his/her goal(s). The term user intent ascription is the attribution of actions to the goal(s) a user will pursue. That is, user intent ascription is the process of determining which actions are attributable to a specific goal or goals. Therefore, for a system to be able to assist the user in pursuing those goals, the system must be capable of ascribing user intent to offer timely, beneficial assistance (Brown, et al., 1998a).

There is a pervasive deficiency in representing the decision-makers in existing systems. The underlying system component that enables decision aids to assist users with their tasks is a user model (Benyon and Murray 1993). An accurate user model is considered necessary for effective ascription of user intent. User modeling is concerned with how to represent the user's knowledge and interaction within a system to adapt the system to the needs, preferences, work flow, goals, skill, expertise, disabilities, etc. of the user as well as the time-criticality, decision-criticality, and uncertainty of the information. Researchers from the fields of artificial intelligence, human-computer interaction, psychology, education, as well as others have investigated ways to construct, maintain, and exploit user models.

The benefit of utilizing a dynamic user model within a mixed-initiative system is to allow that system to adapt its information presentation over time to a specific user. To realize this benefit, the user model must effectively represent the user's knowledge and intent within the system to accurately predict how to adapt the system. The elicitation, specification, design, and maintenance of an accurate user model is necessary for effective ascription of user intent. Research issues include the following.

- Defining the necessary and sufficient methods of eliciting, specifying, designing, and maintaining autonomous and collaborative, semi-autonomous (i.e., human-in-the-loop) user model-based decision aids capable of identifying and assessing information intensive environments.
- Determining the appropriate goals to pursue and tasks to perform
- Performing the tasks given the appropriate level of autonomy.

Of importance in distributed and/or collaborative environments is the ability to model the user's mental

model of who is doing what for/to whom. That is, who is responsible for achieving a goal?

Augmented Visualization

Interactions between the user and system, multiple users, and between systems expand the visualization in temporal sense, just like the increase of field of view in spatial domain. Advances in various modalities in interactions such as speech (to include input ,i.e., speech recognition, and output, e.g., 3D audio), gaze, gestures, virtual reality, and anthropomorphic interfaces not only increase the naturalness of the dialog between humans and computers, but they also augment visualization though bandwidth improvement. However, for effective and efficient communication of the information, determining the best interface modality for information “display” is necessary.

Research to determine what modalities best support overall visualization (in the broadest sense of the word), the requirements for their use, and ways to effectively integrate them into existing environments is an area ripe for future research. One of the most import issues in visualization research is to leverage human capabilities in information perceptualization and interaction skills. These capabilities vary among a user population. Several researchers have chosen a user model-based approach for determining the appropriate modality for information display (Karagiannidis, et al. 1995, Stephanidis, et al. 1996, Horvitz and Lengyel 1997, Brown, et al. 1999). Given a user model of the user's expertise and history of performance on the planning task, one can improve the efficiency by tailoring visualization and interaction tasks to the user. A mixed-initiative system can further explore interactive visualization tailored to particular goals. Interactive visualization should also accommodate temporally varying events such as change of goal-sub-goal relationships through user interactions.

Information Dynamics

We anticipate that new information in any nontrivial domain of interest may arrive asynchronously, and thus, the system must plan in a continuous fashion. That is, one cannot assume complete and consistent information *ab initio*¹. Instead, a knowledge source is only partially complete at any point in time and may contain substantial noise. If new information demands further elaboration, the execution of information-gathering actions may be required to produce additional details. Therefore, steps of a knowledge plan may have to be executed before the final plan is generated. The result is that planning for knowledge goals, like planning for the primary collaborative task itself, is by nature a continual process of planning and execution that must be interleaved. The user(s) needs to be able to visualize how the information “unfolds” over time.

¹ To be sure, we may never have complete and consistent information. Knowing this is the case is just as important.

Although the ideas are extensions and adaptations of existing results in the machine-planning and learning literature, the application of such results to human-information interaction is novel. Results from Cox (1996) provide a foundation for building a theory of errors in information interaction and recovery from it. Further research reported in Veloso, Pollock, & Cox (1998), Cox, & Rasul (unpublished), and Cox & Veloso (1998) contribute a foundational theory of rational-driven sensing monitors that link changes in the environment to adaptation of plans and goals in continuous planning environments. This foundation applies generally, despite the class of planner being used (e.g., state-space or partial order) or the domain being manipulated. Mapping to an information domain will require making the necessary interpolations to algorithms and appropriate development of representations. Note, however, that current research has already demonstrated the fundamental relationship between planning and information acquisition (Cox, 1996; Etzioni & Weld, 1994; Ram & Hunter, 1992).

Common understanding information visualization methods attempt to provide “the right information to the right person at the right time.” User modeling is concerned with how to represent a specific user's knowledge and interaction within a system to adapt the information presentation to the needs, preferences, work flow, goals, skill, expertise, disabilities, etc. of the user. Determining the best of various interaction modalities with the information given the user model further augments the user's understanding of the information presented. The ability to model the dynamic nature of information as well as the time-criticality, decision-criticality, and uncertainty of the information allows users to visualize the “unfolding” of the information over time. These methods are important features of mixed-initiative systems.

Visual Collaboration

The second form of visualization deals with the collaboration between a user and his/her mixed-initiative system (human/system), between users (human/human), and between systems (system/system). There are many examples of environments where multiple users must collaborate to achieve a goal (e.g., air campaign planning, job shop scheduling). Single user environments imply one knowledge plan to help the user achieve his/her goal. However, in an environment where several users work together to achieve a goal, a “higher-level” knowledge plan emerges from the individual knowledge plans of each user. It easy to lose track of who the other users are, who has a critical task that must be accomplished, who controls what information, etc. The ability for individuals to “see the big picture” of the emerging knowledge plan—who are the other users, what are the origins of the information, is the information credible, how does one knowledge plan affect the overall plan, etc.—is vital to effective and efficient collaboration.

Information visualization of the collaborative process can aid users of mixed-initiative systems in understanding the systems' performance. This section investigates several human-centered issues for visual collaboration.

Credibility

As previously mentioned, an advantage of mixed-initiative systems is their ability to perform tasks on the user's behalf. As a result of this advantage, the question arises concerning how to build human confidence in the system's abilities to "do the right thing." Tseng and Fogg present an overview of credibility in computing technology (1999). They contend credibility is synonymous with the term "believability." The phrases "trust in information", "trust in the output", and "accepting the advice" all refer to the credibility of the computer system.

There are two extremes of the credibility issue. At one end, the skeptic refuses to rely on data provided by these systems. Baecker et al. (1995) state one goal of human-computer interaction research is to ensure the user has a feeling of control. Users may feel they have lost control when they have no idea concerning what the system is doing nor what sort of processing the system performs to transform raw data into (possibly) useful information. Klein (1997) argues that "experts prefer to build their own mental models rather than rely on the aggregation and analyses of subordinates who are less skilled." These "subordinates" include computer systems. On the other end of the spectrum is "blind trust," or over reliance on the decision aids. Kilpatrick (1999) states that "as we train more and more reliance on computer systems...we are not training the related common sense, and for lack of a better word, skepticism of computer data. [We] are growing [users] that will blindly trust computer-supplied data, which will make them very vulnerable..." Interactive, visual user models can address the lack of credibility problem. These user models allow the user to collaboratively help the system modify the user model and allow users to build their own mental models of the systems capabilities (Brown 1998).

Information Pedigree

The ability to determine the origin of information lends greatly to the credibility of this information. The process of determining the origins of information is termed *information pedigree*. For a mixed-initiative system to reason about building a knowledge plan for information visualization, it must also reason about the state of information held by other agents including the human user(s). This state may contain irrelevant, incorrect, dissonant, and/or missing information. Building a knowledge plan with uncertain information can result in plan abandonment and therefore excessive replanning, or perhaps worse, achievement of the wrong goals. Therefore, we are concerned with the information's pedigree. That is, we desire to know the answers to questions such as "who 'owns' the data/information?", "are they a trusted agent?",

"how was the information derived?", and "how current is the information?" The answers to these questions may be obtained over time and by many disparate means. Methods to visualize this entire process of information pedigree development are needed and are paramount to increasing system credibility.

Visual Team Planning

The visualization of abstract relationships involved in collaborative planning includes the extraction of pertinent parameters as well as a graphical presentation of the information. Graphical attributes such as locality, proximity, size, and color can be used to represent entities and events occur in a planning process such as goals and precedence relationships. The plan representation and the associated goal-subgoal relationships can be expressed as associated graphs. Interaction tasks are simultaneously defined to facilitate the exploration of the data set and accommodate the dynamics during the planning process. Different levels of details in visualization are created to present the major and minor goal relationships in a networked hierarchical fashion.

A surprising number of domains of interest in the real world involve distributed, collaborative "teams" of individuals interacting to achieve a common goal. While visualization for one user of one mixed-initiative system is important, an equally important goal is presenting a consistent, meaningful representation of the data/information to the team given distributed heterogeneous information sources, applications, and individuals.

A "consistent, meaningful representation" connotes an "intelligent system" where the intelligence is based not only on knowledge of the environment, to include the (possibly dynamic) information sources, but also the specific users of the team. Of utmost importance is the utilization of methods for unifying the intelligent systems, decision support, collaborative technologies and the capabilities and limitations of humans, teams, and work organizations into effective distributive network operating units. Research tasks include teamwork task analysis, tactical and strategic planning perspectives and decomposition, problem-solving strategy engagement, team decision-making in naturalistic environments, and domain knowledge engineering for "chunking" of information displays. Methods include cognitive work analysis, team-in-the-loop simulations, computational modeling, ethnographic studies of the field of practice, and concept prototyping.

User Interface Planning as a Step toward Visualization Planning

In the context of teams of people collaborating in some common task, the goal to produce a tailored information view can become a very complex one to achieve. To

present the proper information that supports an individual in the team and allows that person to visualize the information with respect to the overall activity, a cooperating system must take in to consideration many factors. To illustrate this complexity, we will consider a simple and restricted subset of the problem. That is, we will examine how a system can begin to assist the user by planning the arrangement and configuration of objects within the user's desktop environment.

As anyone knows who has used a modern operating system with a windowing interface, the desktop easily and frequently becomes crowded with multiple overlapping windows of various sizes and content, many of which are no longer relevant to the user's current activity. Clutter can become so intrusive as to interfere with user productivity and to make it difficult to locate the most relevant window for a desired activity. The subsequent sections will examine the problem of planning to reduce window clutter and the associated implementation in the PRODIGY planning architecture.

Planning Clutter-Free Screens

Seeking to cast the problem of information presentation and visualization as a planning task, we examine the reduced goal to achieve *screen clarity* (i.e., to prevent "window clutter"). In the single-user case that ignores the content of windows, clarity can be achieved if a suitable amount of clear space is included and minimal percentage of window overlap is maintained. One solution is to allow only a particular number of windows at a given time. When the number increases beyond the threshold, the system can iconify the least recently used window. However, this solution does not easily scale to more complex solutions that incrementally add further knowledge about the user and the user's task. A knowledge-based solution is to use an automated planner to make the decision as to which window to affect and in what manner.

The Prodigy4.0² system (Veloso, *et al.*, 1995) employs a state-space nonlinear planner and follows a means-ends analysis backward-chaining search procedure that reasons about both multiple goals and multiple alternative operators from its domain theory appropriate for achieving such goals. A domain theory is composed of a hierarchy of object classes and a suite of operators and inference rules that change the state of the objects. A planning problem is represented by an initial state (objects and propositions about the objects) and a set of goal expressions to achieve. Planning decisions consist of choosing a goal from a set of pending goals, choosing an operator (or inference rule) to achieve a particular goal, choosing a variable binding for a given operator, and deciding whether to commit to a

² Note that Prodigy4.0 is the generative planning subsystem in the overall PRODIGY planning and learning architecture.

possible plan ordering and to get a new planning state or to continue subgoaling for unachieved goals. Different choices give rise to different ways of exploring the search space. These choices can be guided by either control rules (see Minton, 1988), by past problem-solving episodes (i.e., cases; see Veloso, 1994), or by domain-independent heuristics (see Veloso and Stone, 1995).

Because the overall task for common understanding visualization is to assist the user to comprehend information relevant to the overall collaborative objectives, an important subtask is to present to the user a tailored picture of that information. To better understand some of the issues involved in this problem, we consider how PRODIGY decides which window to iconify or move when managing a screen layout for windows.

The Window Domain³

This domain explores the problem of alleviating window clutter using a simple world composed of a screen and an arbitrary number of windows. The screen is divided into four area quadrants to render windows and a horizontal rectangle or bar across the bottom to hold icons. When not iconified, a window is contained in exactly one quadrant and fills the entire area. When iconified, the window is moved to the icon bar. Operations on windows are a subset of normal window operations: move a window; minimize a window; and restore a window. A window cannot be resized in this domain.

Several states exist for windows and quadrants. First a window can be part of the set of windows being used by the user for the current task. All windows in this set will be considered "active." All windows not active are considered to belong to a suspended user task. Second a window can be an icon (minimized), or not an icon. By definition an iconified window is never part of the active set. Third a window can be on top of another window (i.e., one window can entirely cover another window in the same area). Therefore a window can be on top (visible on the screen), or not be on top (there is some other window on top of it). Whether a window is on top or not is inferred. Finally, an area can be "clear" or not clear. A clear quadrant indicates that there are no windows currently in the area.


Fig. 1 shows a PRODIGY operator in this domain. The MoveOffOf operator has two window and two area variables. The operator preconditions are that the window to be moved (w) is on top, that the window is not in the quadrant to where it will be moved (a), and that it occupies some other quadrant (a1), and that it is currently on another window (w-below).

³ The full domain and problem definition are available at the following URL.

<http://www.cs.wright.edu/~mcox/mii/adendum/>

The effects of applying the move operator is to assert the following facts:

1. <w> is now in the new quadrant <a> and not in the old one, <a1>.
2. <w-below> is now on top and <w> is no longer on top of it.
3. <a> is not clear.
4. If another window, <w1>, is in the area moved into by <w>, then it is no longer on top and <w> is on top of it and overlaps it.



Example Planning Episode

The example presented here has an initial state consisting of four windows, W1-W4, stacked in quadrant A1 with W1 on the bottom and W4 on top. The goal is to clarify the screen by achieving a clutterfree environment in all quadrants. This can be achieved in each quadrant by not having overlapping windows (i.e., one on top of another). The solution is a simple plan of first moving W4 off of W3 and into A2, W3 to A3, and W2 to A4. However, the simplicity of the problem and solution belies the complexity of the process that is necessary to produce this example. The domain required multiple inference rules, operator constraints, and search control rules to instantiate

the subtle heuristics that managed the decision choices. The firing of one such control rule can be seen in the PRODIGY trace below and its definition in Fig. 2, example variable constraints appear in Fig. 1, and indications of inference rule applications are visible in the plan window of Fig. 3.

Creating objects (W1 W2 W3 W4) of type WINDOW

Creating objects (A1 A2 A3 A4) of type AREA

2 n2 (done)

4 n4 <*finish*>

5 n5 (clutterless a1) [3]

7 n7 <infer-no-clutter a1>

8 n8 not (overlaps w4 w3) [5]

Firing select binding rule SELECT-TOP-WINDOW-BINDING-FOR-MOVEOFFOF at node 9

10 n10 <moveoffof w4 a2 w3> [11]

11 n11 <MOVEOFFOF W4 A2 W3> [5]

12 n12 not (overlaps w3 w2) [4]

Firing select binding rule SELECT-TOP-WINDOW-BINDING-FOR-MOVEOFFOF at node 13

14 n14 <moveoffof w3 a3 w2> [11]

15 n15 <MOVEOFFOF W3 A3 W2> [4]

16 n16 not (overlaps w2 w1) [3]

Firing select binding rule SELECT-TOP-WINDOW-BINDING-FOR-MOVEOFFOF at node 17

18 n18 <moveoffof w2 a4 w1> [11]

18 n19 <MOVEOFFOF W2 A4 W1> [3]

Achieved top-level goals.

#<PRODIGY result: T, 0.11 secs, 19 nodes, 1 sol>

The Prodigy4.0 trace above shows the state of the planner at each decision cycle. Although the system has to only solve the single goal of making area A1 clutterless, it has to manage this without changing the clutter-free state of the remaining areas. That is, it cannot put multiple windows in areas A2, A3, or A4.. This is accomplished by the selective choice of bindings for the variable <a> (i.e., the area to which a window is moved) enforced by control rule SELECT-TOP-WINDOW-BINDING-FOR-MOVEOFFOF.

The definition of this control rule is contained in Fig. 2.

(CONTROL-RULE

SELECT-TOP-WINDOW-BINDING-FOR-MOVEOFFOF

(if (and (current-operator MOVEOFFOF)
(type-of-object <Windo> WINDOW)
(true-in-state (on-top <Windo>))
(true-in-state (clear <QTo>))))

)

(then select bindings

Fig. 2. Example control rule




Fig. 3. PRODIGY goal tree and plan

Extending the Domain to Visualization Planning

The section above illustrates a problem approach rather than a solution. Many issues arise when contemplating an extended view. The domain above assumes a single user which is disengaged from the solution itself. The content of the windows is ignored, and the user-task is unknown. Indeed, no context is included. Furthermore, even though the domain is simple, and only a subset of a realistic situation, the complexity of the search space is such that without heuristic control the generation of plans is impractical. Finally the user may not trust such solutions when produced.

Despite such factors, however, this start represents an preliminary approach from which the idea of planning the visualization can be explored. The implementation allows us to see how far the notion of planning for visualization at the interface can be extended. The next stages of experimentation include the addition of knowledge concerning the content of each window and the relationship windows have to the users tasks. This information can be provided by an external module that sets an initial problem definition file. For example, the

active attributes for windows signal those windows which are being employed by the user for the current problem-solving task.

Another important addition is direct engagement by the user in the process of managing the interface and the visualization. A hook already exists in the PRODIGY architecture that allows any decision made by the system to be made by the user. Taking advantage of this feature, we can allow the user to override decisions the system intends or to help the system make such choices when uncertainty exists.

Many other directions exist from which to expand the implementation. At this time the research has just begun, so few results are available to bias these directions. However, embedding a more complete knowledge of user and window characteristics will enhance mixed-initiative decisions made by future implementations.

Characteristics of the user

- Expertise
- Team relationship
- Cognitive workload and errors
- Personal domain knowledge
- Outstanding information requests

Window Characteristics

- Syntax - how windows are associated hierarchically and linked visually
- Semantics - what functions are provided by the information or tool in each window
- Conceptual Model - how the contents of the windows contribute to overall workflow

Conclusions

One of the most important issues in visualization research is to leverage human capabilities in information perceptualization and interaction skills. These capabilities vary among individual users in a given population and teams of users. Given a model of the user's needs, preferences, work flow, goals, skill, expertise, disabilities, etc. as well as the time-criticality, decision-criticality, and uncertainty of the information and history of performance on the planning task, a mixed-initiative system can improve the efficiency by tailoring visualization and interaction tasks to improve a user's performance in achieving his/her goals. Interactive visualization should also accommodate temporally varying events such as change of goal-sub-goal relationships through user interactions.

Related to the interactive tailoring of the displayed information is the visualization of the collaborative process

between user and system, multiple users, and interacting systems. This visualization can help build confidence in the data abstractions and therefore trust in the analyses the systems provide. The visualization of the information pedigree improves system credibility as well. Knowing where the information originated from, the uncertainty of the information, what pre-processing occurred to transform the data to information are all vital to increasing the level of trust in the information. Finally, the ability to visualize the interaction between the various components in a distributed, collaborative work environment has the benefit of providing a consistent, meaningful representation of the information and tasks involved. This representation, tailored for each member of the team can insure a common understanding of the tasks that are being performed.

Acknowledgements

This paper is supported by the Dayton Area Graduate Studies Institute (DAGSI) under grant #HE-WSU-99-09, by the Information Technology Research Institute, and by Wright State University. We wish to thank Bob Eggleston, Gil Kuperman and Mike McNeese for their technical suggestions, our co-workers for allowing us to flesh out a number of ideas contained herein, and Captain Alex Kilpatrick for his insights into credibility. We also thank Boris Kerkez for helping implement the PRODIGY example from Section 4.

REFERENCES

- Baecker, R.M., et al. (1995), *Readings in Human-Computer Interaction: Toward the Year 2000*, Second Edition, Morgan Kaufmann Publishers, Inc.
- Benyon, D. and Murray D. (1993), "Adaptive systems: from intelligent tutoring to autonomous agents", *Knowledge-Based Systems* 6(4), pp. 197-219.
- Brown, S. M., (1998), *A Decision Theoretic Approach for Interface Agent Development*, Ph.D. Dissertation, Air Force Institute of Technology.
- Brown, S. M., Santos Jr., E., Banks, S. B., & Oxley M. E. (1998a), "Using explicit requirements and metrics for interface user model correction," In K. P. Sycara & M. Woolridge (Eds.), *Proceedings of the Second International Conference on Autonomous Agents* (Agents '98), pp. 1-7.
- Brown, S. M., Santos Jr., E., & Banks, S. B. (1999), "Intelligent Interfaces for Decision-Theoretic Systems," submitted to the *First Asia-Pacific Conference on Intelligent Agent Technology* (IAT 99), Hong Kong, 15-17 December 1999.
- Cox, M. T. (1996). *Introspective multistrategy learning: Constructing a learning strategy under reasoning failure*. (Tech. Rep. No. GIT-CC-96-06). Doctoral dissertation, Georgia Institute of Technology, College of Computing, Atlanta.
- Cox, M. T. & Rasul, G. (unpublished). *Human interaction with automated planners through the manipulation and visualization of goal change*.
- Cox, M. T., & Veloso, M. M. (1999). *Case-based mixed-initiative planning: Supporting combined human and machine decisions*. In preparation.
- Etzioni, O., & Weld, D. (1994). A softbot-based interface to the internet. *Communications of the ACM*.
- Horvitz E., Lengyel J., "Perception, Attention, and Resources: A Decision-Theoretic Approach to Graphics Rendering", *Proc. of the 13th Conference on Uncertainty in Artificial Intelligence*, August 1997.
- Kilpatrick, A. (1999), personal communication.
- Klein, G., *Implications of the Naturalistic Decision Making Framework for Information Dominance*, Technical Report AL/CF-TR-1997-0155, United States Air Force Armstrong Laboratory, July 1997.
- Mitchell, C. M., & Sundstrom, G. A. (1997). Human interaction with complex systems: Design issues and research approaches. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 27(3), 265-273.
- Minton, S. (1988). *Learning search control knowledge: An explanation-based approach*. Boston: Kluwer Academic.
- Narayanan, S., Walchli, S. E., Reddy, N., & Balachandran, R. (1997). Model-based design of an information retrieval system for a university library. *International Journal of Cognitive Ergonomics*, 1(2), 149-167.
- Ram, A., & Hunter, L. (1992). The use of explicit goals for knowledge to guide inference and learning. *Applied Intelligence*, 2(1), 47-73.
- Karagiannidis C., Koumpis A., Stephanidis C., "Media/Modalities Allocation in Intelligent Multimedia User Interfaces: Towards a Theory of Media and Modalities", *Proceedings of the 1st International Workshop on Intelligence and Multimodality in Multimedia Interfaces: Research and Applications*, Edinburgh, UK, 13-14 July 1995.
- Stephanidis C., Karagiannidis C., Koumpis A., "Integrating Media and Modalities in the User-Machine Interface", *Proceedings of the 1st International Conference*

on *Applied Ergonomics*, Istanbul, Turkey, 21–24 May 1996, 256–261.

Tseng, S. and Fogg, B. J. (1999). Credibility and Computing Technology. *Communications of the ACM*, 42(5). 39–44.

Veloso, M. (1994). *Planning and learning by analogical reasoning*. Springer-Verlag.

Veloso, M., Carbonell, J. G., Perez, A., Borrajo, D., Fink, E., & Blythe, J. (1995). Integrating planning and learning: The PRODIGY architecture. *Journal of Theoretical and Experimental Artificial Intelligence*, 7(1), 81–120.

Veloso, M. M., Pollack, M. E., & Cox, M. T. (1998). Rationale-based monitoring for continuous planning in dynamic environments. In R. Simmons, M. Veloso, & S. Smith (Eds.), *Proceedings of the Fourth International Conference on Artificial Intelligence Planning Systems* (pp. 171–179). Menlo Park, AAAI Press.

Veloso, M. M., & Stone, P. (1995). FLECS: Planning with a Flexible Commitment Strategy. *Journal of Artificial Intelligence Research*. 3, 25–52.